Architecture Is Capital Allocation

Why Semiconductor Investment Decisions Begin Before the Capital Meeting

Dana Xiadani
Founder & Chief Architect, Living Cipher
dana@livingcipher.com | www.livingcipher.com
Living Cipher Analysis 004 | August 2026 | Version 1.0
DOI: **10.5281/zenodo.22150344**

Central claim

In capital-intensive semiconductor systems, architecture can preconfigure capital allocation by creating technical dependencies that later implementation, manufacturing, commercial, and policy choices turn into financial commitments. This paper does not propose a new theory of traceability, value-driven design, risk management, staged funding, engineering-change propagation, or technical authority. It proposes a semiconductor-specific operating representation at the commitment boundary: a sparse trace from material premise to dependent commitment, go-forward value basis, ordinary financial-model term, realistic operational/economic recoverability, evidence state, admissibility condition, accountable domain reviewer, and capital decision authority. Its distinctive decision use is selective release: common-path commitments can proceed when they retain positive go-forward value; premise-dependent commitments return to the inherited financial or real-options decision rule rather than a new Analysis 004 score; evidence-generating commitments can be released deliberately when their decision value and sequencing justify the cost; and unresolved irreversible commitments can be held or resized to preserve economically meaningful option value.

Abstract

Semiconductor capital allocation is commonly represented as a financial decision applied to a technical program. That sequence obscures an important dependency: architecture can constrain much of the capital decision before formal authorization occurs, while later implementation, manufacturing, commercial, and policy premises determine how much of that exposure actually crystallizes. Process node, die partitioning, memory hierarchy, package technology, verification scope, supplier set, qualification path, manufacturing geography, and capacity exposure can become financial assumptions as the acceptable implementation set narrows. This paper proposes a semiconductor-specific operating representation—an architecture-to-capital trace linking material premises to the commitments that depend on them, the go-forward value basis for each commitment, the ordinary financial-model term affected, the realistic operational/economic value recoverable if those premises fail, the evidence state of each premise, admissibility conditions, and accountable domain reviewers and capital decision authorities. It builds explicitly on design-to-cost, requirements traceability, risk management, engineering-change propagation, modularity, real-options, Stage-Gate, discovery-driven planning, integrated architecture-cost modeling, value-driven systems engineering, and independent technical-authority practice rather than claiming those mechanisms are new. Public cases from Intel, NVIDIA, and CHIPS implementation show why the additional linkage can matter in semiconductors: node economics can depend on external volume; fabless architectures can create large supply obligations; a policy boundary can strand inventory and contractual commitments after a product is built; and sovereign capital can deliberately alter ordinary milestone conditions. An illustrative demonstrator then decomposes a prospective semiconductor commitment boundary, maps its components into the ordinary economic model—including incremental NRE, cash-flow or working-capital effects, contractual exposure, and expected loss where applicable—and shows how common-path work can proceed while premise-dependent irreversibility remains optional. It does so without replacing the substrate-selection, trust/admission, capital-structure, or expected-value decisions established in prior Living Cipher analyses and ordinary finance. The argument is not that engineers should replace finance. It is that capital diligence is incomplete when the technical and implementation premises that make the financial model possible are detached from their evidence state and from the domains accountable for validating that evidence.

Keywords: semiconductor architecture; capital allocation; technical diligence; chiplets; advanced packaging; manufacturing capacity; capital formation; architecture governance; CHIPS Act

1. Capital Exposure Begins Before the Capital Meeting

Formal semiconductor capital approvals can occur after product definition, architecture selection, initial implementation, supplier conversations, and manufacturing assumptions have already narrowed the option set. The timing varies by firm and program; the claim here is not that finance is always introduced late, but that technical dependency can accumulate before the full economic exposure is formally authorized.

Throughout this paper, architecture means the allocation of system functions, interfaces, state, data movement, and implementation classes that establish consequential technical dependencies. Architecture constrains—but does not uniquely determine—later choices such as process node, package, supplier, manufacturing location, or capacity. Market requirements, portfolio strategy, and policy can precede architecture and can also force it to change. Accordingly, “architecture-to-capital” is shorthand for a trace that begins with architectural constraints and continues through implementation, manufacturing, commercial, and policy premises as economic exposure crystallizes. The title is intentionally provocative: architecture is an upstream determinant of capital allocation, not the first or only capital decision.

Capital allocation is used here in the managerial—not accounting—sense: the commitment of scarce financial resources to a technical path. It can include R&D and NRE, IP and tooling, supplier prepayments, inventory and working capital, capacity reservations, qualification, and property, plant, and equipment. The paper does not imply that each item is capitalized on the balance sheet or treated identically under accounting rules.

A workload has been partitioned. A process node has been assumed. Memory bandwidth has implied a memory technology and package. Interfaces have created IP and verification obligations. A foundry path has narrowed. Test and qualification scope have expanded. Volume forecasts have begun to determine whether the implementation is viable at all.

The formal capital decision may occur later, but architecture has already preconfigured the shape, timing, reversibility, and supplier dependence of much of the eventual commitment.

The scale can be material even for fabless companies that do not own fabs. NVIDIA reported $95.2 billion of manufacturing, supply, and capacity commitments as of January 25, 2026, with substantially all due through fiscal 2027; the company states that these commitments reflect datacenter-scale production and longer ordering horizons across current and future product architectures.[11] The point is not that architecture alone caused that figure. It is that architecture, forecast demand, supplier lead times, and capacity constraints become economically coupled long before finished systems reach customers.

Analysis 001 argued that architecture determines where capital becomes irreversible.[6] The present paper makes a narrower governance and diligence claim: material capital assumptions should remain traceable to the architectural premises that create them, to the evidence supporting those premises, and to the domain competent to validate that evidence.

That broader idea has substantial intellectual ancestry. Design-to-cost and concurrent-engineering methods have long treated cost as an engineering design variable rather than a downstream accounting result.[7] Real-options and modularity research has examined how architecture creates or destroys option value.[8][9] Stage-Gate methods formalized staged commitment and go/kill decisions decades ago.[10] Discovery-Driven Planning goes closer still: it calls for critical assumptions to be made explicit, tested at checkpoints, and funded progressively as knowledge improves.[19] Systems-engineering practice goes further than simple requirement-to-test traceability: NASA guidance calls for assumptions and design rationale to be documented, requirements to remain consistent with architecture, and changes to be assessed for cost, schedule, architecture, design, and interface impact.[14][16] Closer prior art exists at the design-finance boundary: Axelsson explicitly models cost and uncertainty in electronic architecture tradeoffs; Papke and Wang integrate parametric cost estimation with the system-architecture digital thread; and Donelli et al. couple design, manufacturing, and supply-chain variables in early value-driven system trades.[21][22][23] The contribution claimed here is therefore not “architecture affects cost,” “integrate architecture with cost,” “couple design and supply,” “trace assumptions into economics,” or “fund in stages.” It is a deliberately narrow semiconductor operating representation at the commitment boundary: a material technical or implementation premise, the committed resources whose value depends on it, realistic recoverability if it fails, the admissibility conditions that can invalidate it, and the accountable evidence and capital authorities governing the next irreversible commitment.

Materiality rule. A premise is material for this paper only if a credible change in that premise could alter a go/stop/resize decision, the architecture or supplier path, the timing or size of a commitment, or the amount of value recoverable if the program changes course. A premise that cannot plausibly change a decision does not belong in the trace. This rule is intended to keep the representation sparse rather than turn every engineering assumption into an investment-committee artifact.

Prior-art boundary: what is established and what is being added

Existing discipline	Already established	Not claimed as novel here	Increment proposed in this paper
Design-to-cost / concurrent engineering	Cost should enter early design tradeoffs; design choices drive lifecycle economics.[7]	That architecture affects cost.	Trace a material technical premise forward to the specific commitments and unwind exposure that depend on it.
Real options / modularity	Flexibility, modularity, waiting, switching, and staged commitment can have option value.[8][9]	That preserving options under uncertainty has value.	Identify which semiconductor premise creates or destroys the option as tape-out, qualification, supply, and capacity become less reversible.
Stage-Gate / technology readiness	Programs can advance through evidence/maturity gates rather than one all-at-once commitment.[10][14]	That staged funding or readiness gates are new.	Tie the next material commitment to the evidence state of the particular premise on which that commitment depends.

Existing discipline	Already established	Not claimed as novel here	Increment proposed in this paper
Requirements traceability / systems engineering	Requirements, assumptions, design rationale, architecture, verification, and change impacts can already be traced; NASA guidance explicitly includes cost and schedule impacts.[14][16]	That linking technical decisions to cost, schedule, assumptions, or verification is new.	Use a premise-centric decision record for semiconductor commitments: dependent contractual/resource obligations, recoverability/reuse, policy or admissibility conditions, evidence state, and accountable domain/capital authority.
Program risk / assumption management	Risks and assumptions can already be recorded with owners, impacts, mitigations, review status, and decision rationale.[14]	That uncertainty, risk ownership, or mitigation tracking is new.	Make the unit of analysis a material semiconductor premise and explicitly map the commitments whose value depends on it, realistic recoverability, and the evidence condition for the next irreversible commitment.
Discovery-Driven Planning	Critical assumptions should be explicit, tested at checkpoints, and funded progressively as knowledge improves.[19]	That assumption testing or learning-conditioned investment is new.	Bind semiconductor premises to the specific contractual/resource commitments they can strand, realistic recoverability, admissibility conditions, and accountable evidence/capital authorities at the commitment boundary.
Integrated MBSE / value-driven architecture economics	Architecture can already be linked to parametric cost, explicit uncertainty, manufacturing, and supply-chain variables in model-based decision processes.[21][22][23]	That an architecture-cost digital thread or design-supply coupling is new.	At the semiconductor commitment boundary, add only the dependency fields not already explicit: dependent obligations, realistic recoverability, policy/admissibility conditions, evidence state, and accountable evidence/capital authority.
Engineering change / change-propagation analysis	Dependency models and networks can trace how design changes propagate across components, manufacturing, supply chains, cost, schedule, and lifecycle consequences.[26]	That dependency graphs, change-impact analysis, redesign-cost propagation, or production/supply-chain consequences are new.	At the semiconductor commitment boundary, bind a material premise to specific go-forward resource and contractual commitments, their realistic operational/economic recoverability, admissibility conditions, and a release state rather than merely predicting design-change propagation.
Architecture governance / evidence vs. commitment authority	Contemporary architecture practice explicitly distinguishes qualified technical evidence and recommendations from organizational commitment to manufacture or tape out hardware.[27]	That technical evidence and capital/tapeout authority should be separate.	Decompose the next semiconductor commitment boundary by premise, go-forward value, recoverability, admissibility, and RELEASE / RELEASE FOR EVIDENCE / RESIZE / HOLD / RETIRE state.

The paper is therefore a synthesis at a boundary, not a priority claim over integrated architecture-cost models, value-driven systems engineering, or the ingredients from which the trace is assembled. Its incremental value should be judged by whether the proposed premise-to-commitment representation surfaces decision-relevant exposure or missing evidence that the organization’s existing systems-engineering, program-risk, design-to-cost, value-driven design, and financial controls do not already make explicit. If an existing MBSE, PLM, requirements, risk, or investment-control system already records the same dependency, recoverability, evidence-state, admissibility, and authority links, a separate architecture-to-capital artifact adds bureaucracy rather than value.

Capital exposure can begin before formal approval, when architecture narrows what later has to be financed.

2. Technical and Implementation Choices Become Financial Assumptions

The bridge from architecture through implementation to finance is not metaphorical. Technical and implementation choices become variables inside the economic model.

Exhibit 1. Architecture-to-capital translation

Architectural / implementation premise	Immediate technical consequence	Capital consequence	Evidence required before commitment
Process node	Density, performance, power, design rules, mask set	Design NRE, mask/tool cost, wafer economics, node-specific capacity	PPA evidence, PDK maturity, yield assumptions, foundry roadmap
Die partition	Monolithic vs. chiplet boundaries, interfaces, redundancy	Additional die, IP, package/test cost, redesign/requalification scope	Partition benchmarks, interface overhead, yield/cost model
Memory hierarchy	Bandwidth, capacity, locality, data movement	HBM/DDR content, interposer/package demand, supplier exposure	Measured bandwidth demand, utilization, traffic model
Package / interconnect	Thermal density, routing, power delivery, assembly	Substrate/interposer cost, known-good-die/test, package capacity	Thermal/power model, package quote, assembly/test feasibility
Supplier / foundry	Qualification path, process portability, availability	Prepayments, dual-source cost, working capital, stranded redesign	Qualified source set, capacity and lead-time evidence
Volume / lifetime	Amortization and utilization	Break-even, capacity reservations, inventory exposure	Customer demand, design wins, persistence of workload
Software / ISA / toolchain	Compiler/runtime/driver compatibility, portability, ecosystem	Software NRE, time-to-revenue, switching cost, adoption risk	Measured workload portability, software readiness, customer migration evidence

When used to support an investment case, a yield assumption is also a financial assumption. So are thermal limits, bandwidth claims, package-yield assumptions, tape-out schedules, supplier qualification paths, and expected useful life. If a material technical premise is physically wrong, a financial model can be internally rigorous and still describe an impossible or uneconomic business.

Intel provides a public example of how tightly architecture, process technology, manufacturing scale, and capital economics interact. Its 2025 annual report states that leading-edge nodes require substantial capital investment years before return, that disaggregated architectures can extend the use of older nodes, and that Intel 14A requires wafer volumes beyond expected internal product demand to achieve economic efficiency.[3] At year-end 2025, Intel therefore made continued pursuit of 14A and successor leading-edge nodes conditional on obtaining significant external foundry volume. By Q2 2026, Intel had committed to completing 14A development while stating that the scale and pace of manufacturing expansion would be dictated by committed demand from Intel products and external design wins.[13] The sequence illustrates that development commitment and capacity commitment can move at different times as evidence and strategic commitments change.

The implication is not Intel-specific. In capital-intensive hardware, technical assumptions are upstream variables in the financial model, not supporting details appended after the model is built.

3. Architecture Starts a Chain of Capital Precommitment

Capital does not become irreversible all at once. It accumulates in layers.

Early architecture exploration may consume engineering labor while preserving substantial optionality. As the design advances, dependent commitments accumulate: RTL and verification, IP licenses, package design, foundry reservations, masks, qualification, inventory, supplier prepayments, and capacity commitments. The key quantity is not simply money already spent; it is the economic exposure that becomes difficult to recover if a material architectural premise changes.

This paper uses capital precommitment as a working term for downstream economic exposure created by a material architecture-linked or implementation premise before the full program has been realized. The term is descriptive, not claimed as a new financial category. The important point is dependency: which already-committed resources lose value if the premise fails, what value can be realistically recovered or reused, what additional rework or requalification becomes necessary, and what delay or option cost follows from changing course.

Operational exposure ledger

For each material premise, keep four quantities separate: (1) committed resources or contractual obligations whose value depends on the premise; (2) realistic recoverable or reusable value; (3) incremental redesign, re-verification, requalification, or supplier-change cost if the premise fails; and (4) delay or opportunity cost associated with waiting or changing course. The first two describe stranded-commitment exposure; the latter two describe the consequence of unwinding or deferring the choice. The dependency is many-to-many: one commitment can depend on several premises, and one premise can affect several commitments. Exposure therefore cannot be summed mechanically across premise rows without resolving shared dependencies and avoiding double counting. These quantities should not be collapsed into a pseudo-precise expected value when probabilities are not defensible; use observed values, bounded ranges, or explicit scenarios instead.

Recoverability terminology. In this paper, recoverability means operational or economic value that remains usable, transferable, cancelable, creditable, resalable, or applicable to an alternative path if the controlling premise fails. It is not an accounting impairment test or a claim about balance-sheet asset recoverability. Recoverability limits downside; by itself it does not establish that a new commitment has positive go-forward value.

The ledger is a dependency-accounting device, not a universal valuation formula. Its purpose is to prevent double counting, distinguish sunk exposure from go-forward economics, and expose dependencies hidden by ordinary PPA or BOM comparison. A premise with modest direct implementation cost can create large economic exposure if many later commitments depend on it.

Capital precommitment is not a quantity to minimize mechanically. A fixed accelerator, leading-edge node, or dedicated package can rationally create more irreversible exposure if expected volume, performance advantage, and value capture justify the commitment. The objective is to make exposure and recoverability legible relative to the evidence and expected value supporting the decision, while also accounting for the opportunity cost of waiting or preserving flexibility.

Nor should already-sunk cost justify continuation. Sunk expenditure helps describe how irreversible an earlier decision became and can inform post-mortem learning, but the go-forward decision should be based on incremental future economics, recoverable value, remaining risk, and alternatives. The trace is intended to prevent escalation of commitment, not rationalize it.

A decision can therefore remain technically reversible while becoming economically sticky. A programmable accelerator can be changed after RTL with tolerable rework; the same functional assumption can become expensive to unwind after package qualification, HBM commitments, customer firmware freeze, and capacity reservation.

Architecture review should therefore expose not only expected performance and cost, but also dependency: if this premise fails across its credible uncertainty range, what committed work loses value, what must be repeated, and what schedule or capacity consequences follow?

4. Partitioning Changes the Commitment Profile

The chiplet-versus-monolithic choice makes the problem concrete. Disaggregation can improve yield, allow different functions to use different process technologies, increase reuse, and reduce the amount of logic that must be redesigned or requalified when volatile functions change. It can also create die-to-die interfaces, known-good-die requirements, package complexity, assembly cost, test burden, power-delivery constraints, and new supplier dependencies.

The economic coupling is already well established. An Open Compute Project business analysis treats chiplet-versus-monolithic selection explicitly as a cost-and-return problem spanning architecture, die-to-die interfaces, packaging, test, yield, process-node choice, reuse, and time to market.[17] Recent cost-aware research makes that coupling quantitative: CATCH models chiplet decisions across chip size, defect density, test cost, I/O, assembly, substrate, and process choices,[1] while ChipletPart shows that the preferred partition can change with packaging, signaling, floorplan feasibility, and heterogeneous technology assignment.[2]

The lesson is not that chiplets are cheaper. Existing design-economics literature already shows that architecture and modularity can be valued as option-bearing choices.[8][9] The semiconductor-specific point is that the preferred partition is endogenous to package, test, node, yield, supplier, and volume assumptions; changing one premise can alter both the technical optimum and the capital exposure attached to it.

Exhibit 2. A generic partition decision and its capital consequences

Option	Technical attraction	Capital attraction	Capital risk
Monolithic leading-edge SoC	Lowest die-to-die overhead; tight integration	Simpler package and fewer external interfaces	Large die/yield exposure; broader redesign/requalification scope; leading-edge NRE concentrated in one implementation
Disaggregated compute + I/O	Process heterogeneity; reuse; modular evolution	Can keep I/O/control on cheaper nodes and reuse expensive compute	Advanced package/test burden; interface IP; known-good-die and supplier coordination
Programmable bridge / FPGA stage	Maximum learning and algorithmic flexibility	Preserves option value while technical assumptions mature	Higher recurring device cost and potentially worse power/area than hardened implementation
Dedicated fixed accelerator	High specialization when workload is stable	Potential unit-cost/power advantage at sufficient scale	NRE, verification, masks, qualification and volume risk become concentrated

The capital question is therefore not “Which implementation is elegant?” It is “Which implementation produces enough lifetime value to justify the commitments required to make it real, given what is actually known today?”

5. Supply Can Reverse the Architecture

Architecture constrains what must be sourced. Supply, policy, and commercial availability constrain which architectures remain economically available.

A technically superior design can lose because its package capacity is unavailable, its memory supplier is too concentrated, its foundry path requires unacceptable prepayments, its process does not support required analog or I/O functions, or its geographic dependency is inconsistent with customer or sovereign requirements.

Intel provides the IDM case. In its January 2026 annual report, Intel stated that Intel 14A required wafer volume beyond expected internal product demand to achieve economic efficiency and that, absent a significant external foundry customer, it might pause or discontinue pursuit of 14A and successor leading-edge nodes.[3] By Q2 2026, Intel had committed to completing 14A development while stating that the scale and pace of manufacturing expansion would ultimately be dictated by committed demand from Intel products and external design wins.[13] The sequence is more useful than either snapshot alone: a leading-edge node program can move from conditional to committed development as evidence and strategic commitments change, while manufacturing expansion remains demand-gated.

NVIDIA provides the fabless case. Its 2026 Form 10-K reports $95.2 billion of manufacturing, supply, and capacity commitments associated with datacenter-scale production and current and future product architectures.[11] In the first quarter of fiscal 2026, NVIDIA reported $4.6 billion of H20 sales before the April 2025 licensing change and said it was unable to ship an additional $2.5 billion of H20 revenue during the quarter.[15] The new U.S. license requirement covered H20 and other circuits meeting specified memory/interconnect characteristics. NVIDIA recorded a $4.5 billion charge: $1.9 billion for H20 product inventory and $2.6 billion for excess H20 purchase obligations.[15] The company also reported that the charge was lower than initially anticipated because some materials could be reused.[15] The episode does not show that H20 was technically or economically irrational when designed. It shows that product-specific supply commitments can acquire measurable stranded exposure when political admissibility changes after architecture, demand planning, and procurement commitments have formed.

Public capital creates a third case. GAO's August 2026 review reported that Commerce had disbursed $13.1 billion as of April 2026 and that awardees generally received direct-funding disbursements after proof of project milestones such as facility construction, equipment installation or qualification, and production requirements.[20] GAO treated Intel separately in its milestone analysis because Intel's August 2025 amendment created unique terms.[20] Intel's proxy states that the amended commercial CHIPS agreement removed prior project-milestone requirements and accelerated $5.7 billion of remaining disbursements while adding an equity relationship with the U.S. government.[4] Release conditions therefore depend on the objective and bargaining position of the capital provider, including strategic and sovereign objectives that can differ from private capital.

Together, the cases make a narrower claim than “policy drives semiconductors.” Architectural and implementation choices can create commitments; supply and demand determine whether those commitments earn a return; and policy can change either the admissible market or the financing conditions after technical choices have been made.

Case trace. NVIDIA H20: admissibility shock and recoverable exposure

Trace element	Public evidence	Economic consequence	What may be inferred
Commercial premise	H20 generated $4.6B of sales in Q1 FY2026 before the April licensing change; NVIDIA said another $2.5B of H20 revenue could not ship during the quarter.[15]	A product with demonstrated demand and a committed supply chain already had material economic exposure.	Architecture and demand planning can create commitments whose value depends on continued market and policy admissibility.
Admissibility event	On April 9, 2025, the U.S. government required a license for H20 and other circuits reaching specified memory/interconnect characteristics.[15]	The rule changed the economically addressable market after product and supply choices had been made.	Political admissibility can become a design and supply-chain variable, not merely an external risk factor.
Stranded exposure	$4.5B Q1 FY2026 charge: $1.9B H20 inventory provision plus $2.6B excess H20 purchase-obligation charge.[15]	Both owned inventory and contractual supply commitments lost economic value.	The observed charge is a realized exposure measure, not the direct cost of the performance characteristics that triggered the rule.

Trace element	Public evidence	Economic consequence	What may be inferred
Recoverability	NVIDIA said the charge was lower than initially anticipated because certain materials could be reused.[15]	Recoverable/reusable value reduced realized loss.	Any commitment trace should separate gross committed resources from the portion that can realistically be repurposed or recovered.

What the public H20 case proves is deliberately limited. Public disclosures allow the trace to populate observed demand, a policy/admissibility event, a realized inventory and purchase-obligation charge, and evidence of reuse. They do not reveal NVIDIA's internal probability estimates, alternative designs, approval process, supplier-contract flexibility, or what decision would have maximized expected value ex ante. The case therefore demonstrates measurable dependency and recoverability—not that a different governance process would necessarily have avoided the charge.

Architecture constrains the supply chain. Supply and policy can constrain the architecture back.

6. Technical Evidence and Capital Authority Answer Different Questions

The separation itself is not irrational. Different disciplines answer different questions.

Architecture and engineering assess whether the proposed system is physically credible and which technical premises must hold. Manufacturing and supply assess whether it can be fabricated, packaged, tested, qualified, sourced, and scaled at the stated assumptions. Product and commercial teams assess whether customers value it at sufficient volume, price, and timing. Finance assesses whether the resulting return justifies the use of capital relative to alternatives. Boards and investment committees retain formal authority over sufficiently material commitments.

Intel's 2026 proxy provides a straightforward example of formal governance: the board annually discusses and approves the budget and capital-allocation plans and reviews those plans with management.[4] The existence of that authority does not mean the board independently determines whether a package will close thermally or a process will yield as modeled. Those technical premises enter the decision through other domains.

The distinction between technical and programmatic authority is not new. NASA, for example, formalizes an independent Technical Authority process as a check and balance alongside programmatic authority; the program/project manager remains accountable for overall outcome while technical authorities provide independent oversight within their expertise.[18] The semiconductor-capital problem is narrower. A capital authority may legitimately accept a risk without possessing the competence to declare missing technical, manufacturing, commercial, or financial evidence resolved. This paper therefore uses domain evidence responsibility to mean accountability for stating what the evidence within a domain supports, what uncertainty remains, and what would falsify the premise—not monopoly on truth, immunity from challenge, or an absolute veto.

Recent architecture work reinforces the same boundary. Reddi's Architecture 2.0 explicitly separates qualified technical evidence from organizational commitment to manufacture hardware and treats tapeout commitment as a distinct governance decision.[27] Analysis 004 does not claim this evidence/commitment separation as novel. Its narrower concern is the structure of the commitment itself: which specific resources or obligations depend on which unresolved premises, what go-forward value they retain, what value remains realistically recoverable, and therefore which commitments may rationally be released, resized, held, or retired at the current boundary.

Exhibit 3. Domain authority inside a semiconductor capital decision

Domain	Domain responsibility / evidence validated	Decision implication	Failure when absent
Architecture / engineering	Whether the architecture is physically credible; required premises, margins, interfaces, and technical uncertainty	Capital authority may accept technical risk, but should identify which technical premise remains unproven and on what evidence.	Financial model funds a case built on assumptions that are technically false or unmeasured
Manufacturing / supply	Whether the design can be fabricated, packaged, tested, qualified, sourced, and scaled at stated assumptions	Supply/manufacturing evidence constrains cost, schedule, yield, and capacity claims	Program discovers manufacturing or capacity limits after design lock
Product / commercial	Whether customers value the outcome at sufficient price, volume, and timing	Commercial evidence determines whether commitments can plausibly be amortized	Technically successful program lacks economic demand

Domain	Domain responsibility / evidence validated	Decision implication	Failure when absent
Finance	Liquidity, financing cost, return thresholds, portfolio tradeoffs, and downside exposure	Finance determines whether the evidenced opportunity clears the organization's capital hurdle	Technical merit consumes capital without adequate economic return
Board / investment committee	Enterprise-level risk acceptance, strategic fit, and authorization of sufficiently material commitments	Final authority can accept residual risk without converting missing domain evidence into a resolved fact.	No accountable authority owns the total exposure
Public / sovereign authority, when relevant	Legal/political admissibility and strategic value assigned to dependency or domestic capability	Can change market access, subsidy, financing terms, or the effective hurdle rate	Private economics are mistaken for the complete policy objective

The correct structure is therefore not "engineering decides" or "finance decides." Each domain should remain accountable for the evidence behind claims it is competent to validate, while the accountable capital authority decides whether to accept the residual combined risk. A capital authority may explicitly accept residual discretionary risk within its authority; it cannot convert a failed mandatory technical, legal, regulatory, safety, or admissibility condition into a resolved fact. An unrecorded substitution of managerial confidence for missing domain evidence is not risk acceptance; it is loss of traceability.

Domain evidence responsibility should also be separable from authorship of the underlying architecture. Where exposure is material, independent technical review, red-team analysis, supplier evidence, measured prototype results, or another credible challenge path should be able to contradict the sponsoring team. Otherwise technical ownership can become another source of confirmation bias. NASA's independent Technical Authority model is one institutional precedent for separating technical oversight from programmatic sponsorship; this paper does not claim that separation as novel.[18]

7. Condition Irreversible Commitments on the Evidence They Depend On

The governance implication is practical but not novel in the abstract: staged product development, venture financing, real-options reasoning, and Discovery-Driven Planning already preserve optionality by delaying or sequencing commitments while assumptions are tested.[8][10][19] The semiconductor-specific question is narrower: what evidence should retire which technical or implementation uncertainty before a particular irreversible commitment becomes reasonable, and what value remains recoverable if the premise later fails?

The distinction matters because the relevant uncertainties are not generic "project risk." They are architecture-specific: workload stability, PPA, timing, interface behavior, package thermals, yield, test coverage, supplier qualification, volume, and capacity. Each has a different competent validator and a different downstream commitment set.

Conceptual release condition

For each material commitment m, identify the technical and implementation premises on which its value depends, the evidence currently supporting those premises, a credible uncertainty range or scenario set, its go-forward value across the surviving admissible paths, the resources exposed if the premises fail, realistic operational/economic recoverability, the ordinary financial-model term affected, and the decision authority accepting any residual risk. The release condition is therefore a traceable argument, not a universal numeric threshold and not a second expected-value model. The resulting action need not be binary: RELEASE, RELEASE FOR EVIDENCE, RESIZE, HOLD, or RETIRE are distinct states. Common-path work may be RELEASED when it retains positive go-forward value across the surviving paths. Premise-dependent commitments are not required to be attractive in every credible scenario: when decision-grade probabilities or an established real-options model exist, the actual risk-acceptance decision remains in the inherited financial framework. When probabilities are not defensible, Analysis 004 exposes scenario dependence rather than inventing weights, and HOLD or RESIZE preserves option value where delay is economically meaningful. RELEASE FOR EVIDENCE requires that the evidence can change a material later decision, that its expected or scenario-based decision value justifies direct cost, delay, and option cost, and that its sequencing is rational relative to other unresolved premises. RETIRE means updated evidence causes an inherited Analysis 001, 002, or 003 condition to fail.

This avoids turning "milestone financing" into a semantic novelty claim. The proposed discipline is narrower: a schedule milestone is insufficient if the premise supporting the next irreversible commitment remains materially unresolved. The representation is implementation-agnostic. If an existing MBSE, PLM, requirements, program-risk, or financial-control system already links the premise to dependent commitments, recoverability, evidence state, and decision authority, that system should be used rather than creating another review artifact.

Exhibit 4. Evidence-conditioned capital release

Stage	Evidence to retire	Example evidence	Capital that becomes reasonable
Architecture exploration	Workload and partition uncertainty	Measured workload, benchmark, architecture trade study	Prototype engineering; programmable implementation
Implementation proof	PPA, timing, interface, and algorithm assumptions	Synthesis/P&R, hardware prototype, memory/traffic measurement	Larger RTL/DV effort; selected IP commitments
Physical realization	Package, thermal, foundry, test feasibility	Package/thermal model, PDK evidence, DFM/test plan, supplier quotes	Tape-out preparation; package NRE; qualification planning
Commercial proof	Volume, price, customer persistence	Design win, paid pilot, credible forecast, purchase/volume commitment	Mask/tape-out spend; inventory and initial production
Manufacturing proof	Yield, quality, qualification, ramp	Silicon validation, qualification, yield learning, stable supply	Scaled production and working capital
Capacity expansion	Demand exceeds qualified available capacity	Committed demand, utilization evidence, return case	Fab/package expansion, equipment, debt/equity/public incentives

CHIPS implementation shows both the usefulness and the limits of milestone-conditioned release. GAO's August 2026 review states that Commerce generally disburses direct funding after proof of milestone completion and reported $13.1 billion disbursed as of April 2026; Intel was treated separately because of the unique terms created by its amended award.[20] Intel's proxy states that the later agreement removed prior project-milestone requirements and accelerated $5.7 billion of commercial CHIPS funding.[4] The contrast is instructive: release conditions are governance instruments chosen by the capital provider, and sovereign capital can rationally trade technical or commercial milestones against strategic objectives.

A private semiconductor program can use the same logic without pretending every project requires a bureaucratic gate. The useful question is whether the next commitment depends on an assumption whose evidence state is still weak enough that waiting, prototyping, resizing, or preserving an alternative has meaningful option value.

The threshold should also reflect the cost of delay. Waiting for perfect evidence can destroy time-to-market or capacity options just as surely as premature commitment can strand capital. The relevant decision is therefore comparative: commit now, preserve an option, run a targeted experiment, resize the commitment, or abandon—using the expected value and reversibility of each alternative.

8. The Architecture-to-Capital Trace as a Decision-Control Representation

A capital review should be able to trace every material commitment back to the premises that give it economic value, and every material architecture-linked premise forward to the commitments it can strand, resize, delay, or force to be repeated. “Architecture-to-capital” is shorthand: once architecture narrows the implementation set, the trace continues through implementation, manufacturing, commercial, and policy premises that can determine whether the architecture creates value. The trace is a many-to-many dependency graph rather than a stack of independent rows; shared commitments must not be counted repeatedly merely because several premises can invalidate them. At minimum, each material node should name the premise, evidence type and date, credible uncertainty range or scenario set, dependent commitments, go-forward value basis, ordinary financial-model term affected, realistic recoverable value, release state, accountable domain reviewer, capital decision authority, and falsifier.

Exhibit 5. Architecture-to-capital trace for diligence

Step	Question	Failure if skipped
1. Consequential claim	What performance, cost, trust, availability, or market claim must the system satisfy?	Program optimizes a metric that does not determine success
2. Architectural assumption	Which workload, interface, node, package, memory, supplier, and lifetime assumptions support the design?	Hidden dependencies enter capital decisions without being named
3. Evidence state	Which assumptions are measured, simulated, quoted, contractually committed, or merely forecast?	Confidence is inferred from presentation quality rather than evidence
4. Commitment + term map	What go-forward value, NRE, working-capital/cash-flow, contractual, inventory, or expected-loss exposure attaches to each premise and commitment?	Different economic species are mixed or double counted; the cost of being wrong is mismeasured
5. Domain evidence responsibility and decision authority	Who is accountable for validating the evidence, and who is accountable for accepting residual risk and authorizing the commitment?	Decision authority silently substitutes for missing technical, manufacturing, commercial, financial, or policy evidence.
6. Release condition	Which inherited financial rule governs the commitment, what evidence must exist, and which release state applies before the next irreversible tranche moves?	Capital moves on schedule, or Analysis 004 silently becomes a competing valuation model
7. Falsifier	What observation would cause the organization to stop, resize, or change architecture?	Sunk-cost pressure converts a testable thesis into a commitment ritual

This trace is a proposed semiconductor operating representation, not a new systems-engineering theory and not a replacement for DCF, NPV, ROIC, portfolio analysis, value-driven systems engineering, integrated architecture-cost modeling, program management, requirements traceability, design-to-cost, discovery-driven planning, or technical-readiness methods. Its unit of analysis is a material premise or constraint whose credible failure can change the value of committed resources or the decision itself. Its purpose is to keep that premise connected to the financial variables, recoverability, admissibility conditions, and accountable evidence authorities that existing methods consume. If the trace does not create a decision delta—new information that could change sizing, timing, option preservation, architecture, supplier choice, or stop/go judgment—it is documentation overhead. Where existing tooling already preserves those links, the representation adds no incremental value.

9. Demonstrator: The Commitment Trace in Operation

The demonstrator is subordinate by construction. It begins only after the earlier Living Cipher models have done their jobs. Analysis 001 determines whether greater irreversibility has earned technical and risk-adjusted economic justification; Analysis 002 treats required observation and machine evidence as admission constraints and preserves programmability where evidentiary logic remains volatile; Analysis 003 allows financing, ownership, supply, portability, and sovereignty to alter the feasible architecture and economic-capture set.[6][24][25] Analysis 004 does not rescore or override any of those outputs. It asks a narrower question at the next commitment boundary: which material premise is the organization about to capitalize, what committed resources depend on it, and what value remains realistically recoverable if it proves wrong?

Exhibit 6. Model inheritance rule

Prior analysis	Controlling output carried forward	What Analysis 004 is allowed to do	What it may not do
Analysis 001 — From Signal to Substrate	Four sequential gates; the fixed-vs-flexible NPV test; Commitment Ladder; E[L] for downside from failed premises.	Decompose the next commitment boundary and map affected components into the existing economic model—such as ΔNRE, cash-flow/working-capital terms, contractual exposure, or E[L]—while preserving recoverability and avoiding double counting.	Create a second substrate score, a second expected-value model, relax the hurdle H, or treat the gross commitment ledger as a new project-cost or expected-loss term.
Analysis 002 — The Trusted Chip Is Not the Trusted Signal	Claim → required evidence → trust boundary → partition → substrate; required admission evidence is non-substitutable.	Carry required evidence/admission properties into the commitment trace and identify capital exposed to satisfying them.	Trade a failed mandatory admission property for higher NPV, stronger evidence elsewhere, or managerial confidence.
Analysis 003 — Sovereign Clinical Compute	Capital structure <-> compute architecture; ownership, portability, dependency, policy, economic capture, and staged NRE release constrain feasible structures.	Trace heterogeneous semiconductor commitments only inside an admissible capital/supply/ownership structure; decompose the next boundary and show when a changed constraint forces re-entry upstream.	Claim staged NRE release as new, or use a downstream exposure ledger to make an inadmissible supplier, financing structure, or control boundary feasible.

Demonstrator setup. Consider a generalized accelerator program. A stable compute kernel has already cleared the measured-advantage and persistence gates in Analysis 001. Consistent with Analysis 002, the evolving observation-evidence extractor remains programmable while stable admission and enforcement invariants remain protected; the demonstrator therefore does not harden an immature evidence model merely to improve unit economics. Consistent with Analysis 003, the assumed supply, ownership, key-control, and portability structure has already passed the applicable admissibility screen. The candidate now under consideration is a modular fixed-silicon accelerator relative to an FPGA baseline. These are inherited conditions, not outputs produced by Analysis 004.

Analysis 003 already establishes that NRE may be released in stages as evidence matures; staged release is therefore not a contribution of Analysis 004. Analysis 004 changes the unit of control. Rather than treating NRE as a sequential category within a financing or infrastructure arrangement, it decomposes the next semiconductor commitment boundary into premise-dependent obligations - including reusable engineering, production RTL and tapeout, package and qualification work, capacity reservations, prepayments, and inventory - and determines which may proceed independently while others remain optional.

First, apply the existing Analysis 001 economic test. Using purely illustrative inputs - not industry benchmarks - let economically capturable unit volume be 100,000, 220,000, and 350,000 units over three periods; contribution-margin advantage of fixed silicon be $90 per unit; required discount rate be 12%; incremental NRE be $25.0 million; the risk-adjusted downside allowance for risks other than the explicitly stressed volume premise, E[L]_other, be $8.0 million; and the organization's hurdle H be $8.0 million. Volume-conditioned stranded commitment loss is deliberately excluded from E[L]_other in this demonstrator so that the same premise is not charged twice. The discounted contribution-margin advantage is $46.24 million, so the fixed implementation produces an incremental NPV advantage of $13.24 million after Delta NRE and E[L]_other. It clears the $8.0 million hurdle provisionally.

Now stress the material volume premise while holding E[L]_other and the remaining inputs fixed. At 60,000, 120,000, and 180,000 units over the same periods, discounted contribution-margin advantage falls to $24.96 million and the incremental NPV advantage becomes -$8.04 million. Holding timing and other inputs constant, the hurdle is crossed at approximately 0.887 times the base volume profile, or about 594,000 cumulative units. This crossover is conditional on the illustrative E[L]_other allowance. For purposes of the demonstrator, assume the base and low-volume profiles bound the current decision-relevant scenario set; the paper does not claim that these illustrative profiles are empirical forecasts. That scenario set spans the Analysis 001 decision threshold. Volume is material because it can change the substrate decision; Analysis 004 does not invent that materiality. If defensible probabilities later become available, the commitment trace can populate a separate E[L]_volume term, and the Analysis 001 term map must ensure that no volume-conditioned loss appears in both E[L]_other and E[L]_volume.

Exhibit 7. Demonstrator — Analysis 001 decision state

Input / output	Base case	Low-volume stress	Decision meaning
Economically capturable volume Q_t	100k / 220k / 350k	60k / 120k / 180k	Same timing profile; only the material volume premise changes.
ΔCM	$90 / unit	$90 / unit	Inherited from the fixed-vs-flexible comparison.
Discount rate r	12%	12%	Illustrative organizational hurdle input.
ΔNRE	$25.0M	$25.0M	Incremental fixed-silicon implementation cost.
E[L]_other	$8.0M	$8.0M	Illustrative expected-loss allowance excluding the explicitly stressed volume-conditioned commitment loss; prevents double counting.
H	$8.0M	$8.0M	Required economic hurdle for additional irreversibility.
PV of incremental contribution margin	$46.24M	$24.96M	Calculated from the Analysis 001 equation.
ΔNPV after ΔNRE and E[L]	$13.24M	–$8.04M	Base clears H; low-volume stress fails.
Equivalent crossover	~594k cumulative units	—	At the same timing profile, the conclusion changes around 0.887× the base forecast.

The 004 trace begins here. It does not calculate another architecture score and it does not ask the organization to choose between an all-at-once COMMIT and an all-at-once WAIT. It asks a narrower question: while a material premise still spans both sides of the existing Analysis 001 decision threshold, which next commitments remain valuable across that uncertainty, which commitments primarily buy decision-relevant evidence, and which commitments should remain optional because their economics depend materially on the unresolved premise?

Exhibit 8. Premise-dependent commitment ledger and selective release at the next boundary

Next commitment	Go-forward value basis + controlling premise(s)	Gross	Recovery / stranded component if released	Release state
Reusable RTL / DV / interface closure	Common-path work remains useful under both FPGA continuation and fixed-silicon paths; workload/partition must remain valid.	$4.0M	$3.0M / $1.0M	RELEASE. Positive go-forward value across credible substrate outcomes; limited current irreversibility.
Package / thermal feasibility	Evidence-generating work; package/thermal feasibility controls later tapeout and package commitments.	$2.0M	$1.5M / $0.5M	CONDITIONAL RELEASE FOR EVIDENCE. Release now or in parallel only if the evidence-sequencing test shows that package/thermal learning has sufficient decision or time-to-market value relative to other unresolved kill premises.
Production RTL closure + tape-out	Value is largely conditional on fixed silicon proceeding; volume must clear the inherited Analysis 001 crossover and physical feasibility must hold.	$9.0M	$0.5M / $8.5M	HOLD. Full commitment remains materially premise-dependent.
Capacity reservation / prepayment	Value depends on volume, supplier capacity, market/policy admissibility, and the selected production path.	$8.0M	$3.0M / $5.0M	HOLD / RESIZE. A smaller option is justified only if ordinary commercial option economics support it.
Initial materials / inventory	Value depends on volume, market access, and the selected production path.	$4.0M	$2.0M / $2.0M	HOLD. Inventory remains materially premise-dependent.
Total	-	$27.0M	$10.0M / $17.0M	Selective-release result: $4.0M common-path work released now; up to $2.0M is a conditional evidence tranche; $21.0M of premise-dependent production, capacity, and inventory commitments remains optional.

For each commitment m, keep existing and prospective exposure distinct. For an already committed resource or contractual obligation, the stranded component is the committed amount less the portion realistically recoverable or reusable if the controlling premise fails. For a prospective commitment under review, the same subtraction represents potential stranded exposure that would be created if the tranche were released; it is not an existing loss or obligation. Across a premise p, aggregate only commitments whose economic value materially depends on p, resolve shared dependencies so nothing is counted twice, and keep redesign, re-verification, requalification, delay, and lost option value separate. Then map each item to the ordinary economic term it actually affects: incremental NRE, operating or working-capital cash flow, contractual exposure, or E[L] only for probability-weighted loss caused by premise failure. The ledger total is not added to the financial model as a second project-cost or expected-loss measure.

Demonstrator input discipline. The recoverability values in Exhibit 8 are illustrative inputs, not outputs of the framework. In a live program each value requires an explicit basis - common-path reuse, transferable IP, contract cancellation rights, supplier credit, resale value, alternative-product applicability, or another observable recovery mechanism - and should be represented as a range when materially uncertain. A release conclusion that reverses under a modest plausible change in recoverability is itself unresolved.

Analysis 004 does not impose a worst-case capital rule. Common-path work can be RELEASED when it remains admissible and has positive go-forward value across the surviving paths. For premise-dependent commitments, the actual accept/reject decision remains with the inherited Analysis 001 or ordinary financial/real-options framework when defensible probabilities and value estimates exist; Analysis 004 exposes which premise controls the commitment and how much optionality is surrendered. When such probabilities are not defensible, HOLD or RESIZE can preserve meaningful option value without pretending that all credible scenarios are equally likely. A commitment can be RELEASED FOR EVIDENCE only when its principal purpose is to reduce an uncertainty that controls a material later decision, the decision value justifies direct cost, delay, and option cost, and the evidence action is appropriately sequenced. Where multiple unresolved premises can independently kill the same later commitment, a technical experiment should not be funded merely because it is informative if a cheaper, faster, or more decisive upstream test could make that experiment unnecessary—unless parallel execution itself has sufficient time-to-market or option value. RETIRE applies when updated evidence causes a controlling Analysis 001, 002, or 003 condition to fail.

Exhibit 9. Demonstrator — capital released versus optionality preserved

Decision output	Illustrative result	Meaning
Capital released now	$4.0M	Reusable RTL/DV/interface work proceeds because it is common-path work with positive go-forward value across the surviving substrate paths.
Conditional evidence tranche	Up to $2.0M	Package/thermal feasibility work is not automatically released; it proceeds only if the evidence-sequencing test justifies running it now or in parallel.
Recovery / stranded within released tranche	$3.0M / $1.0M	The currently released $4.0M tranche retains substantial reuse value; recoverability limits downside but is not the reason for release.
Capital not yet released	$23.0M	$2.0M is conditional evidence spend and $21.0M is premise-dependent production, capacity, and inventory exposure.
Decision delta	Selective release + evidence sequencing	The program advances common-path work, buys the next most decision-relevant evidence deliberately, and does not treat uncertainty as a reason either to freeze all work or to release all capital.

This is the demonstrator's central result. The unresolved volume premise does not force an indiscriminate wait. Nor does the base-case Analysis 001 result justify releasing the entire $27.0 million boundary. The trace first separates common-path work from premise-dependent work. In the illustrative state above, $4.0 million of RTL/DV/interface work proceeds because it has positive go-forward value under both the FPGA-continuation and fixed-silicon paths; $3.0 million of that tranche is also realistically recoverable or reusable, leaving $1.0 million of current stranded exposure. The $2.0 million package/thermal activity is a conditional evidence tranche rather than an automatic release. It proceeds now or in parallel only if its decision value and time-to-market value survive the evidence-sequencing test. The remaining $21.0 million of production, capacity, and inventory commitments stays optional pending evidence.

Evidence itself is therefore subject to capital allocation. The $2.0 million package/thermal activity is a technical evidence action; it does not resolve commercial volume. If commercial volume can independently kill the fixed-silicon path, the organization should compare the cost, latency, decisiveness, and option value of the available evidence actions before automatically buying both. Suppose a separate commercial evidence action capable of materially narrowing the economically capturable volume range before the held commitments are released costs $1.5 million and imposes an estimated $1.0 million of time-to-market or option cost. Such an action might consist of paid customer validation, independently evidenced design wins, binding or probability-weighted volume commitments, or another test whose acceptance criteria can move the decision-relevant scenario set relative to the approximately 594,000-unit crossover. The all-in cost of that information is therefore $2.5 million. It would be incorrect to justify the action merely because $2.5 million is less than the $17.0 million potential stranded component that would arise if the full boundary were released—or less than some larger worst-case exposure. If defensible probabilities exist, the information should be valued through the ordinary financial model using the probability-weighted decisions it can change. If probabilities are not defensible, the disciplined comparison is scenario-based: can the evidence move the decision-relevant range to one side of the inherited crossover, and is that decision delta worth the direct cost, delay, and option cost? The trace identifies where information has decision value and which evidence should be purchased first without pretending that a stress exposure is an expected loss.

If updated evidence moves the decision-relevant volume range above the approximately 594,000-unit equivalent crossover, the held production and capacity commitments return to Analysis 001 and the ordinary financial model for any actual release decision, subject to the remaining Gate 4 and capital-authority conditions. If evidence moves the range below the crossover, the implementation remains programmable while fixed-silicon commitments are resized, held, or abandoned under Analysis 001. If a mandatory observation or machine-admission condition fails, Analysis 002 forces stop or repartition regardless of economics. If financing, ownership, supplier, portability, or sovereign conditions make the assumed path inadmissible, Analysis 003 forces restructure or repartition upstream. Analysis 004 therefore decomposes and sequences premise-dependent commitments; it does not overrule the models that determine whether the architecture itself is technically, evidentially, economically, or structurally admissible.

The demonstrator produces a concrete decision delta without creating a fourth competing optimizer: continue the commitments that remain useful under uncertainty, deliberately fund evidence that can retire a controlling premise, and withhold or resize the commitments that purchase irreversibility before that premise has earned it. If this separation cannot change timing, sizing, option preservation, architecture, supplier choice, or stop/go judgment, the trace has failed its own materiality test and should be discarded.

The governing question is therefore not merely whether the architecture should be funded. It is: which economic consequences of the architecture have earned the right to become irreversible at this commitment boundary?

10. What Investors Should Ask Before Funding the Architecture

For venture, growth, public-market, strategic, or sovereign capital, the architecture question can be reduced to a small set of diligence tests.

For an external investor, the trace is not itself a valuation model. It matters only when a material premise changes the distribution or timing of cash flows, gross margin, capital intensity, future funding or dilution needs, market access, contractual or inventory exposure, capacity requirements, or the value of preserving an alternative. Those effects should then enter the ordinary valuation and portfolio framework. If no credible premise changes those economics, the trace has little investment value.

- Which architectural assumptions determine the largest portion of the capital requirement?
- Which of those assumptions have been demonstrated on the actual workload, and which remain forecasts?
- Across the credible uncertainty range or scenario set for each critical premise, where does the architecture, capital requirement, supplier path, or economic conclusion change?
- At what volume does the selected implementation become economically superior to the best more-flexible alternative?
- Which supplier, package, process, memory, or geographic dependency is embedded in the design?
- How much of the capital already committed is recoverable if the architecture changes?
- What evidence—not merely what calendar milestone—should justify the next irreversible commitment?
- Who captures the economics if the architecture succeeds, and how capital-intensive is that capture?
- What is already reflected in supplier roadmaps, capacity commitments, subsidies, and market expectations?

The last question matters because architectural truth and investment return are not the same thing. A supplier can be technically indispensable and economically commoditized. A manufacturing layer can receive enormous revenue while requiring still more capital and producing mediocre returns. A small architectural control point can capture disproportionate value if it is scarce and difficult to substitute.

Technical diligence therefore has to continue past “does it work?” into “what must be built if it works, who must finance that build, who captures the resulting economics, and what evidence is still missing?”

11. What Would Falsify the Argument?

This paper does not argue that every architecture decision deserves board attention, that finance lacks legitimate capital authority, or that engineers are naturally superior capital allocators.

The argument weakens where architecture is highly reversible, NRE and supply commitments are immaterial, qualified capacity is abundant and substitutable, or the financial result is insensitive to the technical premises being traced. In those cases, the architecture-to-capital trace adds process without enough decision value.

The governance proposal also fails if preserving domain-specific evidence responsibility does not improve decisions after accounting for slower execution, organizational complexity, confirmation bias, and the tendency of technical teams to defend architectures they helped create. It adds little if ordinary requirements traceability, MBSE/PLM, risk registers, design-to-cost, program controls, and investment reviews already surface the same premise-to-commitment dependencies, recoverability, admissibility conditions, and authority boundaries with equal clarity.

Those are empirical counterweights, not rhetorical caveats. A useful application should surface a material exposure, crossover, missing evidence item, recoverability issue, or option that the ordinary review process would otherwise miss and that can plausibly change a decision. If it does not, the added machinery is unjustified. Technical evidence should remain attributable to competent and, where material, independent review; no domain receives immunity from challenge.

Validation path

This paper is a framework and public-case analysis, not a controlled empirical demonstration that architecture-to-capital tracing improves investment outcomes. That distinction should remain explicit. A retrospective validation would sample both adverse outcomes—disclosed semiconductor write-offs, cancelled or resized programs, supply-commitment charges, and capacity decisions—and matched successful programs that reached production without material stranded exposure. It would reconstruct the material technical and implementation premises and evidence available before each commitment, then test whether the trace discriminates between consequential and non-consequential uncertainty earlier than the ordinary design-to-cost, integrated architecture-cost, technical-readiness, program-risk, and financial-review processes actually used. Successful controls are necessary; otherwise the exercise risks hindsight bias and becomes merely an elegant way to explain known failures after the fact.

A prospective test is stronger: apply the trace to live programs before major commitments and compare decision latency, redesign cost, avoidable write-offs, false-positive stops, and option preservation against comparable programs using existing controls. Four tests can be stated in advance: (1) incremental discovery — does the trace surface a material dependency, recoverability issue, or missing evidence item not already explicit? (2) decision delta — can the discovery change architecture, supplier choice, commitment size or timing, option preservation, or abandonment? (3) realized decision effect — does it actually change a material decision rather than merely improve documentation? and (4) net value — does the avoided exposure or improved decision quality exceed the delay and process cost introduced by the trace? If the method fails these tests, its incremental value is weak.

12. Before the Capital Meeting

Semiconductor capital decisions are often represented as financial approvals applied to technical programs. In practice, market requirements and strategy begin shaping the option set earlier, and architecture then converts many of those intentions into concrete technical dependencies whose economic consequences become progressively harder to reverse.

Architecture can turn selected premises into RTL, IP licenses, packages, foundry paths, qualification programs, supplier dependencies, capacity commitments, and inventory. Manufacturing, commercial, and policy premises can then amplify, reverse, or invalidate those choices. Each step changes both the expected return and the cost of being wrong.

Finance remains necessary because technical merit does not establish economic return. Engineering remains necessary because financial rigor cannot validate physical assumptions it is not equipped to test. Manufacturing, commercial, governance, and sometimes sovereign actors contribute different constraints again.

The useful boundary is therefore not between engineering and finance. It is between assumptions whose evidentiary status is explicit and assumptions whose economic consequences are being embedded in capital decisions without an accountable technical, commercial, manufacturing, or policy basis.

> The capital meeting should authorize commitments and residual risk. It should not be the first place the organization discovers which technical assumptions have already shaped the capital obligation.

For capital-intensive semiconductor systems, the practical discipline is: trace only material architecture-linked and implementation premises to the commitments that depend on them; preserve accountable domain ownership of the evidence supporting those premises; establish the go-forward value basis for each commitment; distinguish existing exposure from potential exposure created only if a new tranche is released; distinguish gross commitment from realistic operational/economic recoverability; map each item into the correct ordinary financial term rather than adding a second cost or loss measure; test crossovers across decision-relevant uncertainty ranges or scenarios; sequence evidence actions when several unresolved premises can independently kill the same later commitment; carry any resulting cash-flow, margin, capital-intensity, market-access, or option-value effect into the ordinary financial model; and make any decision to commit before the evidence is complete explicit, attributable, and economically legible.

Scope and disclosure boundary.

Living Cipher develops private architecture-to-capital decision methods for evaluating technical forcing functions, compute partitioning, supply-chain constraints, evidence maturity, and capital consequence. This paper presents the public analytical thesis and deliberately does not disclose proprietary scoring methods, client-specific diligence procedures, supplier rankings, or protected implementation models.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

During the preparation of this work, the author used ChatGPT (OpenAI) primarily to assist with manuscript organization, structural editing, and refinement of language and presentation. The architectural and analytical framework, technical analysis, interpretation of evidence, judgments, and conclusions are the author’s own. The author reviewed, verified, and edited the manuscript and takes full responsibility for its content.

Selected References

1. A. Graening, J. Talukdar, S. Pal, K. Chakrabarty, and P. Gupta, “CATCH: a Cost Analysis Tool for Co-optimization of Chiplet-based Heterogeneous Systems,” 2025. https://arxiv.org/abs/2503.15753
2. A. Graening, P. Gupta, A. B. Kahng, B. Pramanik, and Z. Wang, “ChipletPart: Scalable Cost-Aware Partitioning for 2.5D Systems,” 2025. https://arxiv.org/abs/2507.19819
3. Intel Corporation, 2025 Annual Report / Form 10-K, filed January 23, 2026. https://www.intc.com/filings-reports/all-sec-filings/content/0000050863-26-000011/intc-20251227.htm
4. Intel Corporation, 2026 Proxy Statement, Board oversight of budget and capital-allocation plans. https://www.intc.com/filings-reports/all-sec-filings/content/0000050863-26-000061/proxy2026.htm

5. U.S. Government Accountability Office, GAO-26-107882, “Semiconductors: Information on Projects Funded to Strengthen U.S. Supply Chain,” December 2025. https://www.gao.gov/products/gao-26-107882
6. Dana Xiadani, “From Signal to Substrate: Architecture, Irreversibility, and the Capital Economics of Custom Silicon,” Living Cipher Analysis 001, Version 1.0, 2026. https://doi.org/10.5281/zenodo.21958390
7. A. Hari, S. Shoval, and J. Kasser, “Conceptual Design to Cost: A New Systems Engineering Tool,” INCOSE International Symposium, 2008. https://www.incose.org/resource/5-1-1-conceptual-design-to-cost-a-new-systems-engineering-tool/
8. D. M. Sharman and A. A. Yassine, “Architectural Valuation using the Design Structure Matrix and Real Options Theory,” Concurrent Engineering: Research and Applications 15(2), 157–173, 2007. https://doi.org/10.1177/1063293X07079320
9. M. Wouters, M. Workum, and P. Hissel, “Assessing the Product Architecture Decision About Product Features — A Real Options Approach,” R&D Management 41(4), 393–409, 2011. https://doi.org/10.1111/j.1467-9310.2011.00652.x
10. R. G. Cooper, “Stage-Gate Systems: A New Tool for Managing New Products,” Business Horizons 33(3), 44–54, 1990. https://doi.org/10.1016/0007-6813(90)90040-I
11. NVIDIA Corporation, Form 10-K for fiscal year ended January 25, 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm
12. Intel Corporation, Form 10-Q for quarter ended September 27, 2025, U.S. Government Agreements / amended CHIPS Direct Funding Agreement. https://www.intc.com/filings-reports/all-sec-filings/content/0000050863-25-000179/intc-20250927.htm
13. Intel Corporation, Form 10-Q for quarter ended June 27, 2026, Intel 14A development and manufacturing expansion disclosures. https://www.sec.gov/Archives/edgar/data/50863/000005086326000157/intc-20260627.htm
14. National Aeronautics and Space Administration, NASA Systems Engineering Handbook, NASA/SP-2016-6105 Rev. 2, requirements traceability and technology readiness guidance. https://www.nasa.gov/wp-content/uploads/2018/09/nasa_systems_engineering_handbook_0.pdf
15. NVIDIA Corporation, Form 10-Q for quarter ended April 27, 2025, H20 revenue, inventory provision, purchase obligations, and recoverability disclosures. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000116/nvda-20250427.htm
16. National Aeronautics and Space Administration, NPR 7123.1A, Systems Engineering Processes and Requirements, requirements traceability, assumptions, decision rationale, and cost/schedule impact guidance. https://nodis3.gsfc.nasa.gov/displayAll.cfm?Internal_ID=N_PR_7123_001A_&page_name=all
17. Open Compute Project, Open Domain-Specific Architecture (ODSA), “Business Analysis of Chiplet-Based Systems and Technology,” November 2023. https://www.opencompute.org/documents/odsa-whitepaper-2024-business-analysis-of-chiplet-based-systems-and-technology-pdf
18. National Aeronautics and Space Administration, Office of the Chief Engineer, “Technical Authority,” independent technical oversight and separation of Programmatic and Technical Authorities. https://www.nasa.gov/technical-authority/
19. R. G. McGrath and I. C. MacMillan, “Discovery-Driven Planning,” Harvard Business Review 73(4), 44–54, July–August 1995. https://hbr.org/1995/07/discovery-driven-planning
20. U.S. Government Accountability Office, GAO-26-109121, “Semiconductors: Commerce Needs Plan to Meet CHIPS for America R&D Requirements,” August 2026. https://www.gao.gov/products/gao-26-109121
21. J. Axelsson, “Cost Models with Explicit Uncertainties for Electronic Architecture Trade-off and Risk Analysis,” INCOSE International Symposium, 2006. https://www.incose.org/resource/cost-models-with-explicit-uncertainties-for-electronic-architecture-trade-off-and-risk-analysis/
22. B. Papke and G. Wang, “Integration of Parametric Cost Estimation with System Architecture — It’s a Dirty Job but Someone Has to Do It,” INCOSE International Symposium, 2018. https://www.incose.org/docs/default-source/events-documents/is2018/is2018---bookabstract.pdf
23. G. Donelli, J. M. G. D. Mello, F. I. K. Odaguil, T. van der Laan, T. Lefebvre, N. Bartoli, L. Boggero, and B. Nagel, “Value-driven Systems Engineering Approach addressing Manufacturing, Supply-chain and Aircraft Design in the Decision-Making Process,” INCOSE International Symposium 33, 463–481, 2023. https://doi.org/10.1002/iis2.13033
24. Dana Xiadani, "The Trusted Chip Is Not the Trusted Signal: How Signal Veracity Changes Hardware Partitioning in Consequential Systems," Living Cipher Analysis 002, Version 1.0, 2026. DOI: 10.5281/zenodo.22019132.
25. Dana Xiadani, "Sovereign Clinical Compute: Capital Formation, Computational Sovereignty, and the Economics of Infrastructure Control," Living Cipher Analysis 003, Version 1.0, 2026. DOI: 10.5281/zenodo.22021064.
26. Arindam Brahma and David C. Wynn, "Concepts of Change Propagation Analysis in Engineering Design," Research in Engineering Design 34, 117-151, 2023. DOI: 10.1007/s00163-022-00395-y.
27. Vijay Janapa Reddi, Architecture 2.0: Principles of AI-Native System and Chip Design, Chapter 11, "What the Architect Owns," August 25, 2026. https://arch2.mlsysbook.ai/book/contents/chapters/11-ownership/

Author Note

Dana Xiadani is Founder & Chief Architect of Living Cipher, where she works at the boundary of computation, architecture, trust, semiconductor strategy, and physical realization. Living Cipher Analysis examines decisions upstream of implementation and capital commitment: what computation deserves to become hardware, which assumptions are being embedded in capital decisions, where technical choices become dependencies, and what evidence should exist before a system becomes expensive to change.

Suggested citation: Xiadani, Dana. 2026. “Architecture Is Capital Allocation: Why Semiconductor Investment Decisions Begin Before the Capital Meeting.” Living Cipher Analysis 004, Version 1.0.